



BLUE GOAT CYBER
Medical Device Cybersecurity

MEDICAL DEVICE CYBERSECURITY GUIDE


IDE Cybersecurity Submission Guide (2026)

How to assemble IDE cybersecurity content: Appendix 3's five recommended elements, the 21 CFR 812.25 requirements behind them, and risk-scaled breadth.

Reviewed 2026-10-22

bluegoatcyber.com

Book a free discovery call
30 minutes with a senior expert
bluegoatcyber.com/discovery-call



IN BRIEF

The cybersecurity content of an IDE application is built against the binding requirements in 21 CFR 812.25, not an eSTAR template. Appendix 3 of the FDA's February 3, 2026 guidance recommends five elements: informed consent disclosure, architecture views, safety scoped security use case views, an SBOM, and cybersecurity labeling, each tied to a specific regulatory subsection, and documentation breadth should scale to the device's actual risk.

The cybersecurity content of an IDE application is assembled against 21 CFR 812.25, not against an eSTAR template. Appendix 3 of the FDA's [February 3, 2026 premarket cybersecurity guidance](#) names five recommended elements: informed consent disclosure, global and multi-patient and updateability architecture views, security use case views for safety-relevant functions, an SBOM, and cybersecurity labeling. Each is cited to a specific 812.25 subsection, which explains its relationship to binding regulations without making the recommended document list independently mandatory. Appendix 4 then scales how much of each you produce to the risk your investigational configuration presents.

This guide is about assembly: what goes in the IDE application, where it lands, and how much of it you owe. If you are still deciding whether cybersecurity applies to your IDE at all, or why Section 524B does not, start with the [IDE cybersecurity requirements guide](#) and come back here to build the package.

Recommended in guidance, required by regulation

Appendix 3 recommendations are nonbinding. Sponsors must satisfy the underlying investigational-plan and informed-consent regulations, but may use an alternative approach that does so.

FDA guidance documents are nonbinding by statute. They state the agency's current thinking and, in the standard disclaimer, describe an approach you may depart from if the alternative satisfies the applicable statutes and regulations. That is genuinely true of the February 3, 2026 guidance.

21 CFR 812.25 is not a guidance document. It is a regulation in the Code of Federal Regulations, and it requires that an IDE application contain a complete investigational plan with, among other things, a description of the device including its significant components and principles of operation, a written protocol, and a risk analysis describing the risks to subjects and how those risks will be minimized. 21 CFR 812.25(g) requires copies of the informed consent materials, which 21 CFR 50.25(a)(2) requires to describe reasonably foreseeable risks to the subject.

Put those two facts together and the structure becomes clear. Appendix 3 is not the source of the obligation. It is the FDA telling you, for a device with connectivity, what a 812.25-compliant investigational plan and risk analysis look like. A sponsor who omits the cybersecurity content is not declining a recommendation; they are filing an investigational plan that does not describe a reasonably foreseeable category of risk to enrolled subjects.

The enforcement path follows the regulation, not the guidance. Under 21 CFR 812.30(b), the FDA may disapprove or withdraw approval of an IDE when the investigational plan is inadequate to protect the rights and welfare of subjects, or when the application contains an untrue statement of material fact or omits

material information. Cybersecurity risk with a credible path to subject harm sits squarely inside that. The 524B refuse-to-accept mechanism is not available, because Section 524B(a) enumerates 510(k), De Novo, PMA, PDP, and HDE and does not reach a submission made under section 520(g). That absence is often mistaken for the absence of any enforceable requirement. It is not.

	Feb 3, 2026 guidance	21 CFR 812.25 / 812.30
Legal character	Nonbinding, current agency thinking	Binding regulation
Creates the obligation	No	Yes
Tells you what compliance looks like	Yes	Only in general terms
Basis for disapproval	Not directly	Yes, via 812.30(b)
Alternative approach permitted	Yes, if it satisfies the regulation	No, the requirement itself is fixed

The practical consequence for a submission: cite the regulation, not the appendix. An investigational plan that says "cybersecurity risks to subjects are analyzed at section 7.4 in accordance with 21 CFR 812.25(b)" reads as compliance. One that says "we have included the Appendix 3 items" reads as a sponsor following a checklist without knowing why.

Appendix 3's five elements, mapped to 812.25

Appendix 3 lists five cybersecurity items as "Recommended" for an IDE submission, each with its regulatory citation attached. Building the package means producing each one and filing it where the regulation puts it.

1. Cybersecurity risks in the informed consent form

Cited to 21 CFR 50.25(a)(2) and 21 CFR 812.25(g).

Subjects enrolling in a study of a connected device must be told that the connectivity carries reasonably foreseeable risks. The relevant obligation is to disclose reasonably foreseeable risks, because 50.25(a)(2) is an unambiguous regulatory requirement about the content of consent.

The disclosure should be specific enough to be meaningful and short enough to be read. Name the connectivity the subject's device will actually use, say in plain language what an attacker could affect, and state what the sponsor and the site will do if a vulnerability is found mid-study. Boilerplate about "data security measures" does not describe a foreseeable risk; it describes a control.

Get this in front of the IRB early. A late change to consent language means re-consenting enrolled subjects, which is a protocol problem, not a documentation problem.

2. Global, multi-patient, and updateability views

Cited to 21 CFR 812.25(c) and (d).

These are architecture views, and they land in the investigational plan's device description and protocol sections because that is what (c) and (d) cover. The global view shows the whole system, including every external interface the investigational configuration exposes. The multi-patient view shows what happens when more than one device, or more than one subject's data, shares an environment such as a study site network or a sponsor-hosted data platform. The updateability view shows the path by which a patched build reaches a device that is already deployed to a site or a subject.

Study builds routinely expose more than production builds do. Debug ports, verbose telemetry, unauthenticated service interfaces left in for troubleshooting, and vendor remote-access tooling all belong on the global view if they exist in the investigational configuration. Documenting the production-intent architecture instead of the study architecture is a substantive misstatement, and 812.30(b) speaks to untrue statements of material fact.

Appendix 3 also treats requirements documentation for these views as recommended. Submit the security requirements each view enforces, not just the diagrams.

3. Security use case views for safety-relevant functionality

Cited to 21 CFR 812.25(c) and (d).

A security use case view traces a single safety-relevant function end to end and shows the security controls along that path. Implant programming is the guidance's own example. The right selection criterion is harm: pick the functions where a security failure produces a physical consequence for the subject, and view those. Functions whose worst case is a data-confidentiality issue generally do not warrant their own view at IDE stage.

Two or three well-chosen use case views beat a comprehensive set of shallow ones, and the selection rationale should be written down.

4. Software bill of materials

Cited to 21 CFR 812.25(c) and (d).

The SBOM is the only sub-element of the broader Cybersecurity Risk Management Report family that Appendix 3 pulls into the recommended column at IDE. Threat model, vulnerability assessment, traceability matrix, security testing, and the cybersecurity management plan are all explicitly downgraded to "could be helpful, but not specifically recommended."

Generate the SBOM from the investigational build rather than from a design document, in a machine-readable format, and regenerate it when the build changes. Include known vulnerabilities in third-party components with a disposition for each. An SBOM with unaddressed critical CVEs in it and no accompanying rationale invites exactly the question you did not want to answer on a study clock.

Because the statutory SBOM obligation in 524B(b)(3) does not attach here, sponsors sometimes skip it. That is a false economy: the investigational stage is the cheapest time to stand up SBOM generation, because the dependency set is smaller and still negotiable.

5. Cybersecurity labeling

Cited to 21 CFR 812.25(f).

Subsection (f) covers the labeling for the investigational device. The cybersecurity content of that labeling is general rather than exhaustive at this stage: what the device connects to, what the associated cybersecurity risks are, and how updates and patches are handled. The audience is the investigator and site staff, and the labeling should tell them what to do about connectivity, not just that connectivity exists.

Where site network configuration matters to the security posture, say so here. An investigator who does not know the device expects a segmented network cannot provide one.

What Appendix 3 does not ask for at IDE

Knowing what to leave out is worth as much as knowing what to include, because over-submission on an IDE creates review questions about documents the FDA did not ask for.

Element	IDE (Appendix 3)	Marketing submission
Informed consent cyber disclosure	Recommended	Not applicable
Global / multi-patient / updateability views	Recommended	Recommended content
Security use case views	Recommended, safety-scoped	Recommended content
SBOM	Recommended	Required, 524B(b)(3)
Cybersecurity labeling	Recommended	Recommended content
Threat model	Could be helpful, not recommended	Recommended content
Security risk assessment (full)	Could be helpful, not recommended	Recommended content
Vulnerability assessment	Could be helpful, not recommended	Recommended content
Traceability matrix	Could be helpful, not recommended	Recommended content
Penetration testing report	Could be helpful, not recommended	Expected
Cybersecurity management plan	Could be helpful, not recommended	Required, 524B(b)(1)

"Could be helpful to submit, but not specifically recommended" is the guidance's own phrasing, and it means what it says: these items are not routinely recommended for submission at IDE stage. The FDA may still request device-specific information needed to evaluate investigational risks.

One nuance that trips teams up. Not submitting a threat model to the FDA is not the same as not having one. You cannot produce a defensible architecture view set, a safety-scoped use case selection, or a 812.25(b) risk analysis without a threat model behind them. Build it; just do not necessarily file it. The same document becomes the spine of the eventual marketing submission.

There is also no eSTAR here. eSTAR is a marketing-submission tool, and reformatting an eSTAR cybersecurity section into an IDE produces a package that answers questions nobody asked and skips the ones 21 CFR 812 actually poses.

Risk-scaled documentation breadth

Appendix 4 sets the rule that determines how much of each element you owe: documentation breadth scales with the risk the device presents.

At the low end sits a device with a single limited interface, such as a wired connection to a clinician workstation with no network path and no multi-device environment. Appendix 4 recommends at least one architecture view for each of three categories: global system, multi-patient harm, and updateability/patchability. Simplicity can reduce the detail and number of views within a category, but it does not turn that recommendation into one total system diagram. Document limited or absent multi-patient and update paths explicitly, and include security use case views for functionality with associated safety risks.

At the high end sits a device that reaches a cloud service, runs on a commercial off-the-shelf operating system, participates in a multi-patient care environment, or supports remote update. The FDA expects the full view set there: global, multi-patient harm, updateability, and security use case views, with the interfaces and trust boundaries actually drawn rather than described in prose.

Two framing points from the guidance carry weight in review. Table 1 in Appendix 4 is a template for organizing content, not a form to be filled in mechanically. And device-specific guidance, where one exists for your product family, takes priority over the general template.

The way to apply the rule is to write the scaling decision down as its own short subsection of the investigational plan. State where the investigational configuration sits on the spectrum, name the factors that put it there, and state what documentation follows. Explain the breadth within each recommended view category and the selection of safety-relevant use cases. This is risk scaling within the recommended evidence set, not permission to replace the three named architecture categories with one total view. An alternative approach must still satisfy applicable regulations.

Revisit the decision if the configuration changes mid-study. Adding a cloud telemetry endpoint to a previously standalone investigational device moves it up the scale, and the supplement is the place to say so.

Assembling the submission

A working sequence that produces the package without duplicated effort:

1. **Build the threat model first**, against the investigational configuration, using production-intent threat, asset, and control identifiers. It is not going in the IDE, but everything that is going in the IDE derives from it.
2. **Write the Appendix 4 scaling rationale**. Decide breadth before producing artifacts, not after.
3. **Produce the architecture views** the rationale calls for, plus the security requirements each view enforces.
4. **Select and draw the security use case views** for functions with a physical-harm path, and record the selection criterion.

5. **Generate the SBOM from the build**, disposition known vulnerabilities, and set up regeneration rather than a one-off snapshot.
6. **Draft the consent language and the investigational device labeling together**, so the risk described to the subject and the risk described to the investigator do not contradict each other.
7. **Write the 812.25(b) risk analysis** so that cybersecurity risks appear as risks to enrolled subjects, with the study's own protections named and the note that those protections disappear at commercialization.
8. **Cross-check citations**. Every cybersecurity artifact in the application should be filed under the 812.25 subsection it answers.

Keep the identifiers stable across all of it. When IDE-stage identifiers diverge from the marketing submission's, traceability breaks, and a broken chain from threat to control to test to patient harm is the leading cybersecurity deficiency pattern the FDA raises.

Practical evidence checklist before filing

Check	Pass looks like
Consent wording	Names the actual connectivity and a foreseeable risk, not only "data is encrypted"
Architecture views	Global, multi-patient, and updateability views exist and match the study build
Use case views	One view for each function with safety risk, such as dosing or programming
SBOM	Generated from the study build, with version identifiers and vulnerability disposition
Labeling	Tells investigators how connectivity and updates work during the study
Internal support	Threat model and risk analysis exist and use the IDs you will reuse at marketing submission

How Blue Goat Cyber approaches IDE cybersecurity submissions

We build the IDE package as the first phase of the premarket engagement rather than as a standalone deliverable, because every artifact in it is one the marketing submission will need again. That means the threat model is built once with production-intent identifiers, the Appendix 4 scaling rationale is written and defensible, the SBOM pipeline runs against the investigational build, and the consent and labeling language is drafted to survive IRB review the first time.

If you have an active study and no cybersecurity documentation behind it, the position is recoverable. The work is reconstructive rather than net new, and it costs far less now than during a marketing-submission review clock.

FAQ

Is FDA cybersecurity guidance legally binding for an IDE?

No, and it does not need to be. The February 3, 2026 guidance is nonbinding. 21 CFR 812.25 is a binding regulation that requires a complete investigational plan and a risk analysis, and 21 CFR 50.25(a)(2) requires consent to describe reasonably foreseeable risks. Appendix 3 tells you how the FDA expects a connected device to satisfy those existing requirements. You may take a different approach if it satisfies the regulation, but you cannot decline the underlying obligation.

Can the FDA disapprove an IDE over cybersecurity?

Yes. 21 CFR 812.30(b) allows disapproval or withdrawal of approval when the investigational plan is inadequate to protect the rights and welfare of subjects, or when the application omits material information. A cybersecurity risk with a credible path to subject harm falls inside both. Section 524B refuse-to-accept mechanisms are not available for an IDE, which is a different point entirely.

Do I need a threat model in the IDE application?

Appendix 3 does not recommend submitting one. You still need to have one, because the architecture views, the use case selection, and the 812.25(b) risk analysis all depend on it. Build it at IDE stage with the identifiers the marketing submission will use.

How many architecture views does my IDE need?

Appendix 4 recommends at least one view for each of the global system, multi-patient harm, and updateability/patchability categories for lower-risk devices, plus security use case views for functionality with associated safety risks. Higher-risk systems may need multiple views in each category. Write down the risk-scaled breadth and any alternative approach; one total system diagram is not the stated recommendation.

Does cybersecurity go in the informed consent form?

Yes, for a connected investigational device. 21 CFR 50.25(a)(2) requires reasonably foreseeable risks to be described, 21 CFR 812.25(g) requires the consent materials in the application, and Appendix 3 names this element explicitly. It should be assessed whenever cybersecurity presents reasonably foreseeable subject risks.

Is penetration testing required before the study starts?

Not by Appendix 3, which places it in the "could be helpful, not specifically recommended" column. The deciding question is whether your threat model identifies a threat whose mitigation cannot be demonstrated any other way. If it does, test it before subjects are enrolled. See [FDA penetration testing scope for 510\(k\)](#) for how scope derives from a threat model.

Where this fits in the cluster

- [IDE Cybersecurity Requirements: The Investigational Device Guide](#) - why 524B does not apply and what the scope actually is.
- [FDA Pathway Cybersecurity Differences](#)

- [De Novo Cybersecurity Submission Guide](#)
- [FDA PMA Cybersecurity Requirements](#)
- [STRIDE Threat Modeling for Medical Devices](#)

Related from Blue Goat Cyber

- [FDA IDE Cybersecurity Requirements](#)
- [FDA Premarket Cybersecurity Services](#)
- [Threat Modeling Services](#)

Sources & primary references

- [Cybersecurity in Medical Devices: Quality Management System Considerations and Content of Premarket Submissions \(FDA, February 3, 2026 final guidance\)](#)
- [21 CFR 812.25, Investigational plan](#)
- [21 CFR 812.30, FDA action on applications](#)
- [21 CFR 50.25, Elements of informed consent](#)
- [Investigational Device Exemption \(IDE\) \(FDA\)](#)


Building an IDE package for a connected device, or repairing one mid-study? [Talk to Blue Goat Cyber about your IDE cybersecurity submission.](#)

Related: the [IDE cybersecurity hub](#) collects every IDE guide, blog post, and pathway comparison we publish.

Current web guide: bluegoatcyber.com/guides/ide-cybersecurity-submission-guide

Talk to a senior medical device cybersecurity expert

Book a free 30-minute discovery call. Scan the code or visit bluegoatcyber.com/discovery-call. You get a written quote within 24 hours.



Informational material. Confirm requirements for your device and market. The manufacturer owns risk decisions, QMS approvals, and regulatory submissions. No regulatory outcome is promised.